

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2010	2. REPORT TYPE		3. DATES COVERED 00-00-2010 to 00-00-2010		
4. TITLE AND SUBTITLE Recognizing faces like humans			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory, Washington, DC, 20375			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 2	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Recognizing faces like humans

Behrooz Kamgar-Parsi, Wallace Lawson, and
Behzad Kamgar-Parsi

A new approach to facial identification makes automated surveillance easier.

Humans are extremely powerful in recognizing faces that they see often. During encounters, we either recognize a face or reject it as unfamiliar. Praise for this ability still echoes in the literature:¹ “The only system that seems to work well under challenging conditions is the human visual system.” While humans do this routinely, a particularly challenging aspect of face-recognition research is the question of rejecting previously unseen faces as unfamiliar. A system with this ability has long been desired:² “The similarity measure used in a face-recognition system should be designed so that humans’ ability to perform face recognition and recall are imitated as closely as possible by the machine.”

The prevailing approach is based on matching and ranking images. Given a test image, a face-recognition algorithm finds its closest match in a database of stored images based on some similarity measure. In ‘closed-world’ applications, where the test image/person is guaranteed to be in the database, if the closest match is found correctly, the test image/person will be identified correctly. By contrast, in ‘open-world’ applications, where the test image/person may not be in the database (as might occur with watchlist surveillance, where we are interested only in recognizing ‘wanted’ subjects), mis-identification may occur regardless of the outcome of the search.

A threshold could be used to decide whether the best match is a correct match. However, establishing the proper threshold value that works well for previously unseen data is very difficult. We have developed an approach that uses an artificial neural network to replicate the human ability to recognize faces. In contrast to most existing approaches, our approach is particularly useful for open-world applications.

Much like humans, our face-recognition algorithm can recognize a certain number of pre-specified ‘target persons’ while rejecting everyone else. In a surveillance application, those individuals may be people on a watch list, and in an access-control application they will be people with authority. Humans




Figure 1. Jennifer Aniston, the leftmost picture, is morphed towards Angelina Jolie, the rightmost picture. The second left picture is still recognizable as Jennifer (positive borderline exemplar, blue dots), while the next picture is not (negative borderline exemplar, red dots). Likewise, morphing Jennifer’s image towards many other people will generate sufficient landmarks to identify and enclose the space belonging to Jennifer (shaded area).

do not match or rank faces/images as a precursor to recognition. Our approach is not based on ranking images, but rather on identifying and enclosing the region in the human face space that belongs to the target person. Consequently, at the test phase, if the image is projected inside that region, it will be identified as the target, but otherwise rejected.

During training, the region is identified with the help of a human critic. We take the image of the target person and morph it towards different facial images from a large database until it becomes borderline acceptable, i.e., significantly different from the target but still recognizable as such (according to the human critic). Next, we morph it even further until it becomes borderline unacceptable, i.e., with some resemblance to the target but not enough to be recognizable as that person.³ Figure 1 illustrates the process.

Typically, a human critic would need to examine morphed images of the target person compared with only some 10 to 20 images from the database to determine the average morphing percentages for the borderline acceptable and unacceptable exemplars. (We have found that different human critics do not

Continued on next page

produce significantly different morphing percentages with regard to a given face.) Applying those percentages and using images in the database, the computer will automatically generate and label a sufficiently large training set. In practice, it is adequate if most of the generated exemplars are projected where intended because the neural net will fit hyperplanes to the generated data. For greater accuracy the database may be divided into subgroups according to gender, race, and age. The human critic may pick slightly different morphed percentages for each subgroup. Next, a three-layer neural network is trained on the two sets of positive and negative exemplars.

Over the past several years, we have collected several thousand facial images at the Naval Research Laboratory (NRL). Furthermore, we have made extensive use of the Face Recognition Grand Challenge dataset⁴ as well as other databases. Extensive algorithmic experiments involving 10 target persons under different indoor lighting conditions and over 10,000 nontarget images have indicated a false accept (false alarm) rate of one in 100,000. The false reject error rate is, however, more difficult to assess. In our experiments such errors appeared to be almost all due to head poses and expressions outside the expected range, e.g., looking far to the side, etc. Of course, in real-life applications of automatic surveillance, cameras will be installed at many different locations, and each is likely to capture more than one image per person. Therefore, it may be of no consequence if the target person is missed in one image or by one camera. That is, the true accept rate does not have to be 100%, but the false accept rate must be extremely low.

The history of face recognition, however, indicates that it is considerably more challenging to obtain a good recognition rate in real-life scenarios, i.e., live systems facing people in person (with many inherent unpredictabilities), rather than algorithms recognizing their pre-recorded facial images. We have been conducting a large-scale system test by taking the system to different buildings (under different lighting conditions) at the NRL. So far, close to 200 people have participated in live experiments. The system has recognized those on its watch list while rejecting all others with no error. The next steps would involve expanded testing, continuation of our effort to further improve the quality of automated morph images, and incorporating age progression.

The support of the US Naval Research Laboratory and the Office of Naval Research is gratefully acknowledged.

Author Information

Behrooz Kamgar-Parsi and Wallace Lawson

US Naval Research Laboratory
Washington, DC

Behrooz Kamgar-Parsi received his PhD in physics from the Catholic University of America. He is currently a senior research scientist at the Navy Center for Artificial Intelligence, where he is engaged in research in computer vision. He is an associate editor of the journal *Pattern Recognition Letters*.

Behzad Kamgar-Parsi

Office of Naval Research
Arlington, VA

References

1. P. Sinha, B. Balas, Y. Ostrovsky, and R. Russel, *Face recognition by humans: nineteen results all computer vision researchers should know about*, **Proc. IEEE** **94**, pp. 1948–1962, 2006.
2. R. Chellappa, C. L. Wilson, and S. Sirohey, *Human and machine recognition of faces: a survey*, **Proc. IEEE** **83**, pp. 705–741, 1995.
3. B. Kamgar-Parsi and B. Kamgar-Parsi, *Methods of facial recognition*, **US Patent pending**, 2007.
4. www.frvt.org/FRGC/ Homepage of the Facial Recognition Grand Challenge project. Last accessed 4 February 2010.